

Towards Subdivision Surfaces C2 Everywhere

Malcolm Sabin^(✉)

Numerical Geometry Ltd., 19 John Amner Close, Ely, Cambs, UK
malcolm.sabin@btinternet.com

Abstract. The conditions for subdivision surfaces which are piecewise polynomial in the regular region to have continuity higher than C1 were identified by Reif [7]. The conditions are ugly and although schemes have been identified and implemented which satisfy them, those schemes have not proved satisfactory from other points of view. This paper explores what can be created using schemes which are not piecewise polynomial in the regular regions, and the picture looks much rosier. The key ideas are (i) use of quasi-interpolation (ii) local evaluation of coefficients in the irregular context. A new method for determining lower bounds on the Hölder continuity of the limit surface is also proposed.

Keywords: Subdivision surfaces · Continuity · Reproduction

1 Motivation

There is a myth among commercial CAD system suppliers that “*Subdivision surfaces are adequate for animation, but they can’t be used for serious CAD because they are not C2 at extraordinary points.*”

Although this myth cannot be completely discounted (like most myths there is some truth hidden underneath), in this raw form it is false.

1. There are subdivision surface schemes which are C2 at extraordinary vertices (see [9, 18, 19] and several others). They may all have other problems, but they are C2.
2. Lack of C2 at isolated points does not matter. The ideas of [18] could be applied after a few hundred iterations of a standard method. This would make the limit C2 everywhere, but it would be totally indistinguishable from the raw limit surface. Also, there is a bivariate interpolation technique which minimises the bending energy of the surface. This is widely regarded as good, but the second derivatives are unbounded at the data sites.
3. The ‘raw’ forms of popular subdivision surfaces [2, 4, 5] are not merely not C2. They actually have unbounded curvature at extraordinary points, and because the unboundedness accelerates at different rates for positive and negative Gaussian curvatures, this can force regions of the limit surface close to an extraordinary point to have the wrong sense of curvature [15, 16]. Such extreme behaviour can be eliminated easily [23, 24] by tuning of the coefficients in extraordinary regions, but this does not cure the lack of C2, and is not widely cited.

The truth underneath is that the raw forms of the popular schemes, and even tuned versions to a lesser extent, suffer from local distortion of the surface over non-infinitesimal regions. It is these artifacts which are visible and objectionable in the reflection line plots, and it is these which need to be eliminated. The basic cause of these turns out to be the same as the cause of our inability to tune these schemes to perfection: the inability to achieve *quadratic generation* in any simple uniform stationary scheme which gives bi-polynomial pieces in the regular regions. This was understood nearly twenty years ago [7,9,10] and that understanding has deepened since [17].

This paper explores the idea of using schemes with fractal limit surfaces, the analog of the quasi-spline curves of [20,21] to give the kind of behaviour which serious CAD needs. [25] has already explored the quad grid case, and so the focus here is on surfaces where the regular part of the mesh contains triangles.

Any full coverage of this will require many more years of work, and so this paper merely sets out the territory and provides some initial results and some ideas for directions which need to be explored. The exploration here is of surfaces with a quasi-interpolating degree of three.

2 Definitions

A scheme is said to **generate** polynomials of degree d if for any bivariate polynomial of degree d we can find a control polyhedron for which that function is the limit.

A scheme is said to **reproduce** polynomials of degree d if for any bivariate function of degree d we can find a control polyhedron for which that function is the limit, and for which the control points all lie on the function. The limit has to interpolate them.

The distinction between these two was not emphasised in the early papers, but it turns out to be fairly important. Over regular grids, most subdivision schemes can generate polynomials of low degree. Those which are B-spline based do not reproduce them for degrees above 1.

A scheme can reproduce polynomials without being an interpolating scheme for arbitrary data. Such a scheme is called a **quasi-interpolating scheme** of degree d where d is the highest degree of polynomial which it can reproduce.

A scheme is called a **stable generator** of degree d if it generates polynomials of degree d and if the maximum change of d^{th} derivative with respect to addition of perturbations of maximum amplitude ϵ is continuous with respect to ϵ with Hölder exponent greater than 1, 0. The issue of stability is critical. The four-point scheme is a quasi-interpolant of degree 3, but not a stable one, because for any sequence of six data points which do not all lie on a cubic, the second derivatives at dyadic points within the central span diverge at a rate depending on the fourth differences of the local control points.

3 Quadratic and Cubic Reproduction in Functional Subdivision over a Regular Triangulation of the Domain

This section exemplifies an approach. It does not preclude further work on higher quasi-interpolation degrees, on different regular connectivities or on higher arities.

3.1 Stencil Sizes

A quasi-interpolating functional subdivision scheme has the property that if the values at all the old vertices lie on a polynomial, then new vertices will also get vertices on that polynomial. If this property holds at every step, then the limit surface will reproduce that polynomial, which we take as a condition for continuity of the relevant degree, and for avoiding certain artifacts.

In order for a new vertex to have that property, the value has to be chosen to match all polynomials of the target degree, d , which means that unless the scheme is actually interpolating, there is a number n of conditions to be fulfilled equal to the number of independent polynomials of that degree. These are tabulated in Table 1.

Table 1. Stencil sizes

Degree	Number
0	1
1	3
2	6
3	10
4	15
n	$(n + 1)(n + 2)/2$

Table 2. Number of available vertices

No of rings	Number of vertices in V-vertex stencil	Number of vertices in E-vertex stencil
0	1	
1	7	4
	13	8
2	19	10
	25	14
	31	18
r	$3(r + 1)(r + 2) + 1$	

We get the new value by taking a linear combination of the values at the vertices in the stencil of the new point, and in order to be able to satisfy all these conditions, we need to have at least as many independent coefficients in those linear combinations. This leads to the unwelcome conclusion that high degrees of polynomial reproduction will require very large stencils, involving a number of rings varying at least linearly with the degree. See Table 2. Quadratics and cubics remain (just) within sensible reach.

We can have more coefficients (i.e. more old vertices in the footprint of the stencil) than conditions to satisfy, and in this case the choice of the values of those coefficients is underdetermined. This can be demanded by symmetry, since the number of vertices in a rotationally symmetric stencil goes up with the number of rings included.

However, in the case of the completely regular domain, symmetry also provides a way of resolving the extra freedom.

3.2 Construction of Cubic Quasi-interpolant

We know that each of the stencils must contain at least 10 entries for cubic quasi-interpolation. There are three e-vertex stencils as well as the v-vertex stencil, and so the total number of entries in the mask must be at least 40. We need some understanding in order to be able to construct such a beast.

This understanding comes from two sources: the first is that the v-vertex stencils must have 6-fold rotational symmetry, and that the e-vertex stencils must have two mirror symmetry axes. The second is that we cannot have good behaviour for general data if we do not have good behaviour for extruded data, in which each mesh direction has the same value at all points along each edge in that direction.

If the data is extruded, then the shape of a cross-section limit curve will be given by a univariate refinement scheme whose mask is given by the row-sums of the bivariate scheme. Further, factorisation of the mask is preserved under the taking of row-sums. Because there are three directions for the row sums, the desired symmetries come out in the wash.

We know a univariate scheme with cubic quasi-interpolation. This was described by Hormann and Sabin [20]. Its mask is

$$2 \left(\frac{[1, 1]}{2} \right)^6 \frac{[-3, 10, -3]}{4}$$

To map this to a triangular grid we need to take the $[1, 1]/2$ factors in pairs, replacing each pair by

$$S = \begin{bmatrix} 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 \end{bmatrix} / 8$$

and finding a kernel which, when multiplied by the cube of this, gives the required row sums.

This is easily found to be

$$K = \begin{bmatrix} -3 & -3 \\ -3 & 26 & -3 \\ -3 & -3 \end{bmatrix} / 8$$

which has the correct row sums to match $[-3, 10, -3]/4$, giving a final mask¹ of $4S^3K$.

$$\left(\begin{array}{cccc|c} & -3 & -12 & -18 & \\ & -12 & -28 & -24 & \\ -18 & -24 & 66 & 144 & \\ -12 & -24 & 144 & 468 & \\ -3 & -28 & 66 & 468 & 754 \\ & & & & sym \end{array} \right) /1024$$

From this we can extract the stencils

$$\left(\begin{array}{cccc} -3 & -18 & -3 & \\ -18 & 66 & 66 & -18 \\ -3 & 66 & 754 & 66 & -3 \\ -18 & 66 & 66 & -18 \\ -3 & -18 & -3 & \end{array} \right) /1024$$

and

$$\left(\begin{array}{ccc} -12 & -12 & \\ 24 & 144 & 24 \\ -28 & 468 & * & 468 & -28 \\ 24 & 144 & 24 \\ -12 & -12 & \end{array} \right) /1024$$

There are enough entries in these stencils (19 and 14 respectively) to give cubic precision, which requires 10 points, but not quite quartic, which requires 15. In any case there is no reason why these coefficients should satisfy the quartic quasi-interpolation conditions. Three of the four cubics and all lower degrees are satisfied because cubic extruded data is matched for three extrusion directions (by symmetry). The fourth cubic condition is based on a linearly varying quadratic, and this is satisfied by mirror symmetry. We therefore assert that this set of stencils satisfies all of the constant, linear, quadratic and cubic quasi-interpolation conditions.

The question remains whether the resulting scheme is stable in the sense that the eigenvalues behave.

¹ K here is normalised to sum to 1: the kernel is thus $4K$. This factor of 4 is analogous to the factor of 2 in the expression for the univariate scheme mentioned above.

3.3 The Basis Function

The basis function is the effect on the surface when one control point is moved. If we start the refinement process itself with a single unit valued control point and all the rest with value zero (cardinal data), successive refinements converge towards the basis function. We can get tessellations of the basis function by convolving each of these with the unit row stencil (the row (left) eigenvector of the component with unit eigenvalue and column (right) eigenvector of all 1s) to get the vertices at limit points. These are shown in Fig. 1. The last of these can be compared with the basis functions of the Butterfly [5] and Loop [4] schemes in regular regions in Fig. 2. Our scheme has significantly smaller wriggles than Butterfly, and is narrower than Loop. In all of these figures, the basis function is shown as far out as the first zero entries. Note that these are not control polyhedra, as the scheme is only quasi-interpolating and the basis function is not a polynomial.

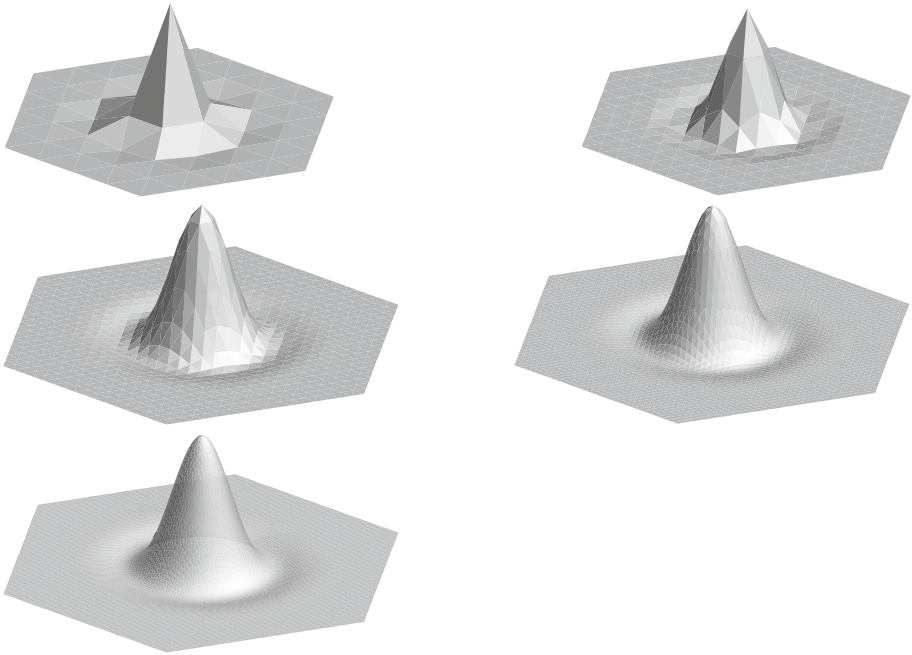


Fig. 1. The first four iterations, in which there is refinement, but no alteration of existing limit points.

3.4 Continuity Analysis

We can address the determination of lower bounds on the Hölder continuity by three routes. The first approach is to apply the bivariate approach of Cavaretta et al. [6], which is capable of showing that whatever the data, a norm of the first

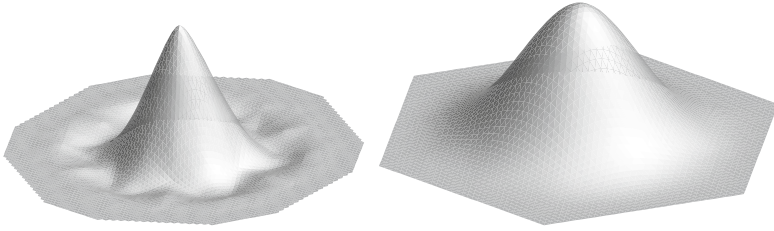


Fig. 2. The basis functions of the butterfly scheme and of Loop in regular regions. (also drawn at four iterations.)

difference of the second or third divided difference contracts with each refinement. This is capable of putting a lower bound on the continuity of the limit surfaces. Unfortunately I have not been able to formulate this approach and have not therefore been able to apply this test.

The second is that of Kobbelt [11], which applies difference operators to the powers of the refinement operation.

I suggest here that a simpler approach may be to check whether the contractivity is true for the basis function. Any other data has its limit surface expressed at any point as a finite weighted sum of basis functions, and so if the basis function is continuous in derivative to a certain degree, then so must any other limit surface. This can be done numerically and is covered in Sect. 3.5 below.

We can also determine upper bounds by checking the magnitudes of the eigenvalues, which can put upper bounds on the level of continuity, by measuring the actual continuity at specific mark points of the surface. This can also be done numerically and is covered in Sect. 3.6 below. Sharper bounds can in principle be determined by eigenanalysis at more mark points. Here only the triangle vertices and triangle centres are considered.

3.5 Contractivity of Differences of the Basis Function

Because we have tessellations of the basis function at different levels, we can, merely by convolving with appropriate difference stencils, determine how the various differences vary from level to level.

Because it is easiest to implement, the tables below are derived by taking the appropriate differences (up to 4th) over the entire domain and then extracting norms. We show here the max-norm and the average absolute value.

Because of the symmetry of the basis, we need only take the first differences in one direction giving Table 3. Similarly the three second differences are just symmetric versions of a single one giving Table 4. We can get the four 3rd differences just by evaluating two of them giving Table 5. The same is true of the 4th differences giving Table 6.

In each table are five columns. The second is the sum of the absolute values of the relevant difference, taken over the tessellations of the basis function shown in Fig. 1. The third is the ratio between each such value and the one above. The fourth column is the maximum absolute value of the relevant difference, and again the fifth holds the ratio between successive refinements.

Table 3. First differences: Because the entries in the second column are sums, the ratios need to be divided by 4 to get the actual ratio of norms. (There are four times as many triangles and therefore four times as many difference entries in each row compared with the previous.) With this correction, both ratios are converging nicely to $1/2$, which, because it is less than one indicates that the scheme is convergent, with a continuous limit.

Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	1.954658364122		0.6252769189585	
1	4.4209983756427	2.2617754881314	0.3706062878079	0.5927074494053
2	9.457042987369	2.1391193083155	0.2077661161766	0.5606114170526
3	19.150115219234	2.0249580386608	0.1054197019496	0.5073960272712
4	38.430854767130	2.0068210727280	0.0529147334585	0.5019434932932

Table 4. Differences of first divided differences: The ratios in the third column now need to be divided by 4, but also multiplied by 2, because the raw figures are differences rather than divided differences. Those in the fifth column now need to be multiplied by 2. Again, the ratios are nicely converging to $1/2$ which is well below 1, and so we can deduce that the scheme is C1.

Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	1.954658364122		0.625276918958	
1	5.195647796224	2.658084855947	0.509341262301	0.260578150970
2	5.915246892456	1.138500361158	0.148838512720	0.292217661784
3	6.111079357843	1.033106389124	0.039238461919	0.263631107314
4	6.162914058208	1.008482085950	0.009987309486	0.254528567085

3.6 Eigenvalue Analysis

This is very straightforward, as we can work solely with the kernel, which is very simple.

$$4K = \begin{pmatrix} -3 & -3 & \\ -3 & 26 & -3 \\ -3 & -3 & \end{pmatrix} / 2$$

Table 5. Differences of second divided differences: We have two different third differences in use here. The ratios in the third column now need to be divided by 4, but also multiplied by 4, because the raw figures are differences rather than divided differences. Those in the fifth column now need to be multiplied by 4. Yet again, the ratios are nicely converging to 1/2 which is well below 1, and so we can deduce that the scheme is C2.

Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	6.3472772348861		1.7968990864628	
1	6.7026946541919	1.0559952568878	0.4211556586767	0.2343791378434
2	4.2440378797963	0.6331838310942	0.0783173770441	0.1859582684705
3	2.2718369140307	0.5353008098362	0.0103408828255	0.1320381659330
4	1.1575101905374	0.5095040860497	0.0013108955912	0.1267682472874
Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	8.457483383656		2.175810559215	
1	19.80725893292	2.341980236248	1.479495404617	0.679974365576
2	18.91360654972	0.954882581874	0.365260442678	0.246881768972
3	11.64307651131	0.615592614804	0.053058514549	0.145262142707
4	6.249221667750	0.536732852496	0.006830399904	0.128733342097

Table 6. Differences of third divided differences: We have three different fourth differences in use here. The ratios in the third column now need to be divided by 4, but also multiplied by 8, because the raw figures are differences rather than divided differences. Those in the fifth column now need to be multiplied by 8. It looks as though we can infer C3.

Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	12.30937074944		3.593798172925	
1	9.351603172048	0.759714152932	0.786811211287	0.218935837080
2	3.573727352791	0.382151304652	0.086012788579	0.109318204095
3	1.044397941094	0.292243318527	0.008115334957	0.094350329662
4	0.280368542187	0.268449918518	0.000710776028	0.087584311946
Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	16.91248082729		4.351621118431	
1	37.84281382575	2.23756728608	2.95899080923	0.679974365576
2	24.47335784142	0.646710837997	0.511077643651	0.172720253829
3	7.817348127387	0.319422785301	0.051984063931	0.101714611423
4	2.094546402655	0.267935669299	0.004790014051	0.092143893521
Level	$\Sigma d $	Ratio	$ d _{max}$	Ratio
0	10.70181694529		2.314131729739	
1	5.6719938630861	0.5300028856857	0.4288936385916	0.1853367434013
2	1.8838035407895	0.3321236916438	0.0439635901700	0.1025046450080
3	0.5026102536695	0.2668060882075	0.0040813827780	0.0928355205346
4	0.1320672373545	0.2627627200007	0.0003559242029	0.0872067684662

The stencils are

$$\begin{pmatrix} 0 & 0 \\ 0 & 26 & 0 \\ 0 & 0 \end{pmatrix} / 2 \quad \text{and} \quad \begin{pmatrix} 0 \\ -3 & * & -3 \\ 0 \end{pmatrix} / 2$$

The subdivision matrix at vertices for the zero frequency Fourier component becomes

$$\begin{bmatrix} 26/2 \\ -3/2 & -3/2 \\ 0 & -6/2 & 0 \\ 0 & 26/2 & 0 \end{bmatrix}$$

whose only non-zero eigenvalues are $26/2$ and $-3/2$. Those of higher frequency Fourier components are never larger.

We may therefore determine the Hölder continuity at vertices to be

$$6 - \log_2(13) = 6 - 3.7004 = 2.2996,$$

because each S factor increases the Holder continuity by 2.

The subdivision matrix at triangle centres for the zero frequency Fourier component becomes

$$\begin{bmatrix} -6/2 \\ 26/2 & 0 \\ -3/2 & -3/2 & 0 \\ -3/2 & 0 & -3/2 \end{bmatrix}$$

whose only nonzero eigenvalue is -3 . Those of higher frequency Fourier components are never larger.

The Hölder continuity here is $6 - \log_2(3) = 4.41$

In fact the continuity might be less than 2.2996, because this analysis only tells us about the continuity at the dyadic points, giving upper bounds on the level of continuity. Other mark points might have different (worse) behaviour. There are an infinite number of them and this is why the joint spectral radius calculation is such a lot of work. However, getting two upper bounds on the right side of 2 means that the method has not failed this test.

4 Triangulations with Extraordinary Vertices

The handling of extraordinary vertices requires amazingly little additional analysis. If a scheme has quadratic or cubic precision, it will have linear precision. We can therefore choose a desired natural configuration, for example by taking that of any C2 scheme, and the coefficients computed to give the reproduction properties will automatically make that natural configuration a self-fulfilling prophecy.

We have three types of stencil to determine:

- (i) That of the EV itself.

Here use of symmetry resolves the underdeterminedness.

- (ii) Those of new vertices² lying within the support of the EV: i.e. whose stencils contain the EV in their interior, so that the topology of the stencil is affected asymmetrically by the presence of the EV.

This is where most of the work lies, in that the vertices which need to be taken into account have to be chosen. The EV spoils the regularity of the topology. Calculating the coefficients to give the required reproduction properties applies as in the regular case.

- (iii) Those of vertices lying further out, in what we can term the ‘*far field*’. Here the topology of the stencil is regular, but the geometry is influenced by the natural configuration. Thus the actual coefficients will differ from the regular case. This effect propagates unboundedly, although hopefully reducing with distance. However, we can choose to resolve the underdeterminedness by minimising the difference from the regular case. Because of the regress of regularity around the extraordinary vertices, if the stencils converge fast enough towards the regular case, we do not have to worry about continuity at any points other than the extraordinary vertices themselves. This means that eigenanalysis at the EVs is a sufficient tool. Thus most of the really hard work is done in the regular case.

In fact at large distances from an EV we have two possible strategies

- (iii.i) to be pure, and insist that all stencils are determined correctly.

In fact, if we are careful, the perturbation of the coefficients from the regular case reduces fast enough with distance that after a certain distance it can be approximated accurately enough (to machine precision) by some simple form, which can be combined when more than one EV is present in the control polyhedron.

- (iii.ii) just to revert to the regular stencil for all new vertices whose support does not touch the EV.

This avoids a lot of implementation problems, and will not spoil the continuity at the EV, because only control points within, or on the boundary of, the support of the EV influence the shape of the limit surface in the immediate neighbourhood of the EV. It will not spoil the continuity anywhere else either, because everywhere outside that neighbourhood has its continuity governed by the regular case.

However, we can expect something to go wrong, and an obvious symptom to be expected will be the appearance of artifacts in the region between the two regimes.

A practical resolution might well be to use (iii.ii) to avoid overlapping influences from more than one EV. In regions where more than one EV has a significant effect, who can say what the desired surface should be? As refinement

² In the case of schemes having multiple kinds of e- and f-vertices (for example, that of [9] or of higher arity) each kind will need individual calculations of this type.

proceeds, the threshold between the two regimes can move in more slowly than the refinement rate, and the artifacts should therefore be spread out and diluted.

4.1 Conformal Characteristic Map

In [13] David Levin conjectured that a stationary subdivision scheme giving the characteristic map

$$u + iv = (s + it)^{6/n}$$

would require a EV mask of unbounded size. Here we have a non-stationary scheme, because the stencils depend on the local layout of control points in the domain, and the corresponding conjecture is that the influence of the presence of the extraordinary point extends unboundedly, even though the EV itself is only accessed by the stencils of a finite number of nearby new vertices.

In the far field we find that the rate of convergence to the regular case turns out to be significantly slower if we enforce cubic reproduction than if we only enforce quadratic, despite the fact that what we are trying to be close to does have cubic precision when the geometry is regular. The Tables 7, 8, 9 and 10 in Annex 2 report $E1 = \sum_i \delta_i^2$ and $E2 = \max_i |\delta_i|$ for reproduction of degrees 2 and 3, for v-vertices along rays from the EV of valencies 5 and 7.

4.2 Other Characteristic Maps

The choice of the conformal characteristic map is entirely an aesthetic one. It would be equally possible to choose as arbitrary natural configuration that of Loop or of the Butterfly scheme. Anything, in fact, which can be shown to be 1:1 and have its characteristic map join smoothly enough to its scaled instance under the next ring.

4.3 Moving Least Squares

Since presentation of the paper I have seen the recent work of Ivriissimtzis. This uses a combination of a standard subdivision scheme (in this Loop, Butterfly or even the degree 1 box-spline) to determine approximate new vertex positions, which are then refined by projection on to a local least squares fit to some subset of the old vertices. This work is following a line which can be traced back through the work of Boyé, Guennebaud and Schlick [22], Levin [14] and McLain [1]. It looks to be almost equivalent to the above ideas, in that it gives a quasi-interpolant (if enough local points lie on a polynomial, the new point will lie on that polynomial), but goes, apparently, more directly to the new points (projecting rather than first computing coefficients).

5 Directions Still Needing Exploration

5.1 Stronger Proof of Degree of Continuity

The analysis of the regular case above is unsatisfactory, with different approaches giving different opinions. The method of chapter seven in Cavaretta et al. [6] needs to be made to work.

5.2 End and Edge Conditions

Before these ideas can be applied to practical situations, the detail of what to do at the edges of the domain of interest needs to be articulated. Unfortunately we do not yet even have end conditions adequately explored for univariate quasi-splines, and this must obviously come first.

5.3 Extension to Higher Degrees

In principle exactly similar constructions can be made, choosing coefficients to give any desired quasi-interpolation degree. Because the stencils for higher degrees become large quadratically with degree, it is unlikely that higher degrees will be of more than academic interest, but the academic interest is there. It seems likely that the extension in this direction will not expose new problems significantly different from those already seen.

5.4 Extension to Solids and Higher Dimensions

Although regular triangular grids do not have completely regular analogues in higher dimensions, the ideas of using subdivision basis functions which reproduce low degree polynomials without being interpolating for general data must be relevant to IsoGeometric Analysis, overcoming the current disadvantage of subdivision that the lack of polynomial generation gives excessive stiffness because the artifacts contribute spurious energy to the solution.

Acknowledgements. My thanks go to a very diligent referee who bothered to construct a counterexample disproving a conjecture in my first draft of this paper and also pointed out many places where the original text was not clear. I also thank colleagues: Cedric Gerot for helping me understand that counterexample, Leif Kobbelt and Ulrich Reif for prompt replies to my emails requesting clarifications, and to Ioannis Ivriissintzis for bringing the moving least squares ideas to my attention.

Annex 1: Computation of Stencil Nearest to a Regular Stencil

In the topologically regular but geometrically irregular case, I suggest that the nearest solution to the regular one should be chosen. This gives some measure of continuity with respect to changing layouts in the abscissa plane. The metric for ‘nearest’ might be chosen by more sophisticated arguments later³, but for the moment I just use the euclidean distance in coefficient space.

Call the number of coefficients c and the number of quasi-interpolation conditions n .

³ For example, to get better continuity when the set of support points needs to change because of changes in the set of local neighbours.

The quasi-interpolating conditions define a linear subspace of dimension $c - n$ in the space of dimension $c - 1$ of sets of coefficients: nearness to the regular coefficients defines a complementary subspace orthogonal to it, and the solution can be found in that subspace by solving a linear system of size $n \times n$.

Let the set of coefficients be a_i , $i \in 1..c$, and let the quasi-interpolation conditions be

$$\forall j \in 1..n \quad \sum_i a_i f_j(v_i) = 0$$

If the regular coefficients are \bar{a}_i , then we can set up the system

$$\forall i \quad a_i = \bar{a}_i + \sum_{k \in 1..n} \delta_k f_k(v_i)$$

and solve for the δ_j .

$$\begin{aligned} \forall j \in 1..n \quad 0 &= \sum_i a_i f_j(v_i) \\ &= \sum_i (\bar{a}_i + \sum_k \delta_k f_k(v_i)) f_j(v_i) \\ &= \sum_i \bar{a}_i f_j(v_i) + \sum_i \sum_k \delta_k f_k(v_i) f_j(v_i) \\ &= \sum_i \bar{a}_i f_j(v_i) + \sum_k \sum_i \delta_k f_k(v_i) f_j(v_i) \\ &= \sum_i \bar{a}_i f_j(v_i) + \sum_k \delta_k \sum_i f_k(v_i) f_j(v_i) \end{aligned}$$

$$\begin{bmatrix} \ddots & & & & \\ & \sum_i f_k(v_i) f_j(v_i) & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \ddots \end{bmatrix} \begin{bmatrix} \vdots \\ \delta_k \\ \vdots \end{bmatrix} = - \begin{bmatrix} \vdots \\ \sum_i \bar{a}_i f_j(v_i) \\ \vdots \end{bmatrix}$$

The actual coefficients can then be determined from these δ_j values.

$$\forall i \quad a_i := \bar{a}_i + \sum_k \delta_k f_k(v_i)$$

The rate of convergence with distance from the EV in an interesting natural configuration can be measured by how $\sum_j \delta_j^2$ varies with distance.

Annex 2: Convergence for Conformal Characteristic Map

These figures in Tables 7, 8, 9 and 10 ($E1 = \sum_i \delta_i^2$ and $E2 = \max_i |\delta_i|$) are disappointing in that the convergence is so slow. A conjecture that the rate of convergence would be d^{-6} has been soundly disproved by these calculations.

Table 7. Valency = 5, reproduction degree = 2

Distance	E1	E2
2	0.0266273481	0.00859447462
4	0.0029902673	0.00098588428
8	0.0004904519	0.00014386747
16	0.000117032	0.00003354449

Table 8. Valency = 5, reproduction degree = 3

Distance	E1	E2
2	0.043294222	0.0178461598
4	0.035778240	0.0131572241
8	0.020705507	0.0069547056
16	0.01054870	0.0035061301

Table 9. Valency = 7, reproduction degree = 2

Distance	E1	E2
2	0.03776611323	0.0112682406
4	0.00295127833	0.0010064401
8	0.00022928609	0.0000913945
16	0.0000595678	0.0000175042

Table 10. Valency = 7, reproduction degree = 3

Distance	E1	E2
2	0.1065766380	0.0439985659
4	0.0327818713	0.0132975787
8	0.0142165239	0.0051765249
16	0.0074573353	0.0025183365

References

1. McLain, D.H.: Two dimensional interpolation from random data. *Comput. J.* **19**(2), 178–181 (1976)
2. Catmull, E.E., Clark, J.: Recursively generated B-spline surfaces on topological meshes. *CAD* **10**(6), 350–355 (1978)
3. Doo, D.W.H., Sabin, M.A.: Behaviour of recursive division surfaces near extraordinary points. *CAD* **10**(6), 356–360 (1978)
4. Loop, C.T.: Smooth subdivision surfaces based on triangles. M.S. Mathematics thesis, University of Utah (1987)
5. Dyn, N., Levin, D., Gregory, J.: A butterfly subdivision scheme for surface interpolation with tension control. *ACM ToG* **9**(2), 160–169 (1990)
6. Cavaretta, A.S., Dahmen, W., Micchelli, C.A.: Stationary Subdivision. *Memoirs of the American Mathematical Society*, no. 453. AMS, Providence (1991)
7. Reif, U.: A degree estimate for subdivision surfaces of higher regularity. In: *Proceedings of AMS*, vol. 124, pp. 2167–2175 (1996)
8. Reif, U.: TURBS - topologically unrestricted rational B-splines. *Constr. Approx.* **14**(1), 57–77 (1998)
9. Prautzsch, H.: Smoothness of subdivision surfaces at extraordinary points. *Adv. Comput. Math.* **9**(3–4), 377–389 (1998)
10. Prautzsch, H., Reif, U.: Degree estimates for C^k -piecewise polynomial subdivision surfaces. *Adv. Comput. Math.* **10**(2), 209–217 (1999)
11. Kobbelt, L.: $\sqrt{3}$ -subdivision. In: *Proceedings of SIGGRAPH 2000*, pp. 103–112 (2000)
12. Velho, L.: Quasi 4–8 subdivision. *CAGD* **18**(4), 345–357 (2001)
13. Levin, D.: Presentation at Dagstuhl Workshop (2000)
14. Levin, D.: Mesh-independent surface interpolation. In: Brunnett, G., Hamann, B., Müller, H., Linsen, L. (eds.) *Geometric Modeling for Scientific Visualization*, pp. 37–49. Springer, Heidelberg (2004)
15. Peters, J., Reif, U.: Shape Characterization of subdivision surfaces - basic principles. *CAGD* **21**, 585–599 (2004)
16. Karciauskas, K., Peters, J., Reif, U.: Shape Characterization of subdivision surfaces - case studies. *CAGD* **21**, 601–614 (2004)
17. Levin, A.: The importance of polynomial reproduction in piecewise-uniform subdivision. In: Martin, R., Bez, H., Sabin, M. (eds.) *IMA 2005. LNCS*, vol. 3604, pp. 272–307. Springer, Heidelberg (2005). doi:[10.1007/11537908_17](https://doi.org/10.1007/11537908_17). ISBN 3-540-28225-4

18. Levin, A.: Modified subdivision surfaces with continuous curvature. *ACM (TOG)* **25**(3), 1035–1040 (2006)
19. Zulti, A., Levin, A., Levin, D., Taicher, M.: C2 subdivision over triangulations with one extraordinary point. *CAGD* **23**(2), 157–178 (2006)
20. Hormann, K., Sabin, M.A.: A family of subdivision schemes with cubic precision. *CAGD* **25**(1), 41–52 (2008)
21. Dyn, N., Hormann, K., Sabin, M.A., Shen, Z.: Polynomial reproduction by symmetric subdivision schemes. *J. Approx. Theor.* **155**(1), 28–42 (2008)
22. Boyé, S., Guennebaud, G., Schlick, C.: Least squares subdivision surfaces. *Comput. Graph. Forum* **29**(7), 2021–2028 (2010)
23. Augsdörfer, U.H., Dodgson, N.A., Sabin, M.A.: Artifact analysis on B-splines, box-splines and other surfaces defined by quadrilateral polyhedra. *CAGD* **28**(3), 177–197 (2010)
24. Augsdörfer, U.H., Dodgson, N.A., Sabin, M.A.: Artifact analysis on triangular box-splines and subdivision surfaces defined by triangular polyhedra. *CAGD* **28**(3), 198–211 (2010)
25. Deng, C., Hormann, K.: Pseudo-spline subdivision surfaces. *Comput. Graph. Forum* **33**(5), 227–236 (2014)